

AMENDMENTS TO THE CLAIMS

Please amend the claims as indicated in the following listing of all claims:

1. (Previously Presented) In a scheduler for computer code, which is encoded in one or more computer readable media, wherein certain operations are likely to stall execution of the computer code and thereby provide latency for completion of one or more pre-executable operations, a method of scheduling certain of the operations, the method comprising:

for one or more sequences of operations that follow a speculation boundary and that define respective dependency chains, including pre-executable operations, which lead to likely stalls, representing speculative copies thereof as duplicate chains; and

scheduling operations of the computer code, wherein the scheduling of operations from the duplicate chains is performed without regard to dependence of respective original operations on the speculation boundary, thereby scheduling certain of the operations above the speculation boundary into position preceding at least one of the operations likely to stall execution of the computer code, wherein the speculation boundary corresponds to an operation upon which the respective original operations depend.

2. (Original) A method, as recited in claim 1, wherein the likely stalls include likely cache misses.

3. (Original) A method, as recited in claim 1, wherein the dependency chains include address chains leading to memory access operations likely to miss in a cache.

4. (Original) A method, as recited in claim 1, wherein the pre-executable operations include prefetch instructions.

5. (Original) A method, as recited in claim 1, wherein the pre-executable operations include speculative operations.

6. (Original) A method, as recited in claim 1, wherein the operations likely to stall execution include memory access instructions.

7. (Original) A method, as recited in claim 1, wherein the operations likely to stall execution include operations selected from the set of:

- a load operation;
- first use of a load operation;
- a store operation;
- a branch operation;
- a multi-cycle computational operation;
- an iterative or recursive operation;
- a communications operation;
- an input/output (I/O) operation;
- a synchronization operation; and
- a co-processor operation.

8. (Original) A method, as recited in claim 1, wherein the speculation boundary is defined by one of:

- a store operation;
- a branch operation;
- a join operation;
- an iterative or recursive operation;
- a communications operation;
- an input/output (I/O) operation;
- a synchronization operation; and
- a co-processor operation.

9. (Original) A method, as recited in claim 1, further comprising: inserting the pre-executable operations into the computer code.

10. (Original) A method, as recited in claim 1, further comprising: profiling the computer code to identify the likely stalls.

11. (Original) A method, as recited in claim 1, further comprising:
upon reaching the speculation boundary, deleting unscheduled operations of the duplicate chains and continuing to schedule respective original operations.
12. (Original) A method, as recited in claim 1, further comprising:
deleting from the original operations, pre-executable operations for which a respective speculative copy is scheduled.
13. (Previously Presented) A computer implemented method of hiding latency in computer code wherein certain operations thereof are likely to stall execution, the method comprising:
identifying sequences of operations that define respective original dependency chains that lead to likely stalls and for at least some of the identified sequences, representing duplicate dependency chains thereof; and
scheduling at least some operations from the duplicate dependency chains above at least one of the likely-to-stall operations,
wherein the operations scheduled from the duplicate dependency chains are scheduled above a speculation boundary that corresponds to an operation upon which the original dependency chains depend and that precedes the at least some of the identified sequences.
14. (Cancelled)
15. (Original) The method of claim 13,
wherein dependency chains are defined at least in part by address dependencies.
16. (Original) The method of claim 13, further comprising:
upon reaching a corresponding speculation boundary, removing otherwise unscheduled operations of the duplicate dependency chains.
17. (Original) The method of claim 13, wherein the operations likely to stall execution include operations selected from the set of:

- a load operation;
- first use of a load operation;
- a store operation;
- a branch operation;
- a multi-cycle computational operation;
- an iterative or recursive operation;
- a communications operation;
- an input/output (I/O) operation;
- a synchronization operation; and
- a co-processor operation.

18. (Original) The method of claim 14, wherein the speculation boundary is defined by one of:

- a store operation;
- a branch operation;
- a join operation;
- an iterative or recursive operation;
- a communications operation;
- an input/output (I/O) operation;
- a synchronization operation; and
- a co-processor operation.

19. (Original) The method of claim 14, wherein the speculation boundary is defined by an operation that has irreversible side-effects.

20. (Original) The method of claim 13,
wherein the operations likely to stall execution include memory access operations.

21. (Original) The method of claim 13, further comprising:
for at least load-type ones of the operations, inserting corresponding prefetch operations.

22. (Original) The method of claim 13, further comprising:
converting load-type ones of the scheduled operations to speculative counterpart
operations.
23. (Original) The method of claim 13, further comprising:
converting load-type ones of the scheduled operations to non-faulting loads.
24. (Original) The method of claim 13, further comprising
responsive to the scheduling of a prefetch operation from one of the duplicate
dependency chains, disposing of a corresponding prefetch operation from a
corresponding one of the original dependency chains.
25. (Original) The method of claim 13, further comprising:
selecting for the scheduling, particular ones of the operations from the duplicate
dependency chains based at least in part on chain length.
26. (Original) The method of claim 13,
wherein the likely to stall operations include memory operations predicted to miss in a
cache.
27. (Original) The method of claim 13,
wherein the likely to stall operations include store-type operations predicted to miss in a
cache.
28. (Original) The method of claim 13,
wherein the likely to stall operations include operations that stall an execution pipeline.
29. (Original) The method of claim 13,
wherein the dependency chains include load-type and prefetch operations.

30. (Original) The method of claim 13,
wherein the dependency chains include operations other than load-type and prefetch
operations.
31. (Original) The method of claim 13,
wherein the dependency chains include operations involved in address calculations.
32. (Original) The method of claim 13,
wherein the duplicate dependency chains are represented as copies of the respective
original dependency chains with speculation boundary dependencies removed or
ignored.
33. (Original) The method of claim 13,
wherein the dependency chains are represented a directed acyclic graph of dependencies
amongst the corresponding operations.
34. (Original) The method of claim 33, wherein the dependencies include one or more
of:
register dependencies;
branch dependencies; and
memory dependencies.
35. (Original) The method of claim 13, realized in an optimizing compiler.
36. (Original) The method of claim 13, realized in a just-in-time (JIT) compiler.
37. (Previously Presented) A computer implemented method of making a computer
program product that encodes program code for which memory access latency is at least partially
hidden on execution thereof, the method comprising:
for operations that form addressing chains that lead to a likely cache miss, representing
speculative copies thereof; and

scheduling the speculative copies without regard to a corresponding speculation boundary, wherein operations of the speculative copies are scheduled above the corresponding speculation boundary and above a preceding operation that is likely to stall,

wherein the speculation boundary, which precedes the operations that form addressing chains, corresponds to an operation upon which the operations that form addressing chains depend.

38. (Original) The method of claim 37,
encoding the scheduled operations as part of the program code.

39. (Original) The method of claim 37,
wherein preceding operation that is likely to stall is a likely cache miss.

40. – 46. (Cancelled)

47. (Previously Presented) An apparatus comprising:
a code preparation facility for transforming schedulable code into scheduled code; and
means for scheduling speculative copies of operations that form dependency chains that lead to a likely stall, the scheduling placing the speculative operations above a preceding at least one other operation that is itself likely to stall, thereby hiding in the scheduled code latency of the speculative operations,
wherein the scheduling also places the speculative operations above a speculation boundary that precedes the operations that form the dependency chains,
wherein the speculation boundary corresponds to an operation upon which the operations that form the dependency chains are dependent.

48. (Previously Presented) The apparatus of claim 47, further comprising:
means for inserting pre-executable operations into the schedulable code, wherein at least some of the pre-executable operations are scheduled by the scheduling means as the speculative operations for which latency is hidden.

49. (Original) The apparatus of claim 47, further comprising:
means for identifying likely-to-stall operations of schedulable code.

50. (Cancelled)